Forging a Future in Memory:

New Technologies, New Markets, New Applications

Ed Doller

V.P. Chief Memory Systems Architect

Non-Volatile Memory Seminar
Hot Chips Conference
August 22, 2010
Memorial Auditorium
Stanford University

Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH MICRON™ PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN MICRON'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, MICRON ASSUMES NO LIABILITY WHATSOEVER, AND MICRON DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF MICRON PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Micron products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Micron may make changes to specifications and product descriptions at any time, without notice.

Micron may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Designers must not rely on the absence or characteristics of any features or instructions marked “reserved” or “undefined.” Micron reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Contact your local Micron sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Micron literature may be obtained by visiting Micron's website at http://www.Micron.com.

*Other names and brands may be claimed as the property of others.

Technology Lifecycle

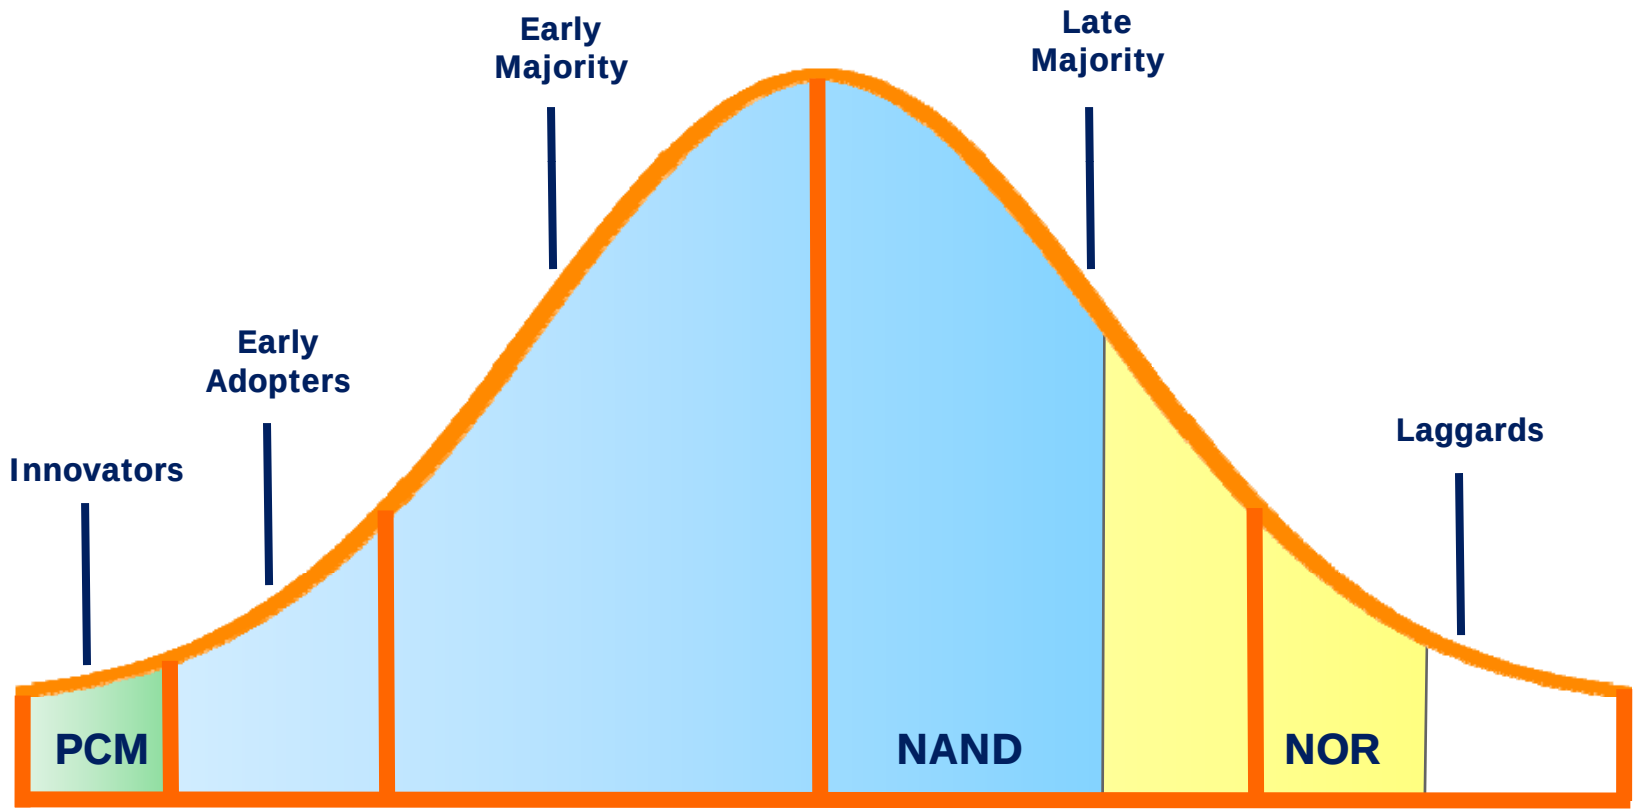

NOR Scaling

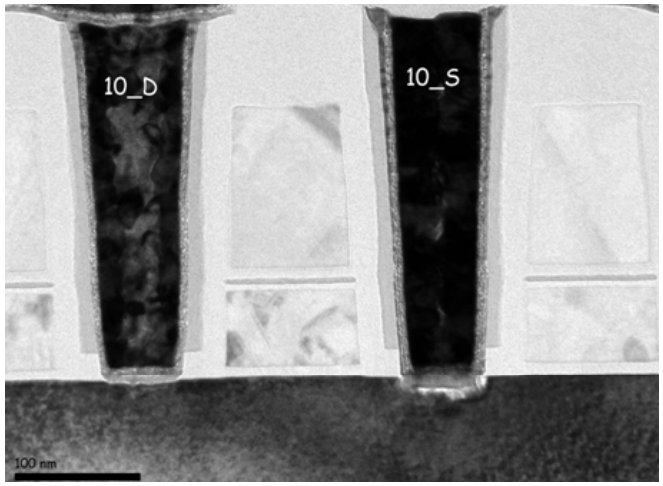

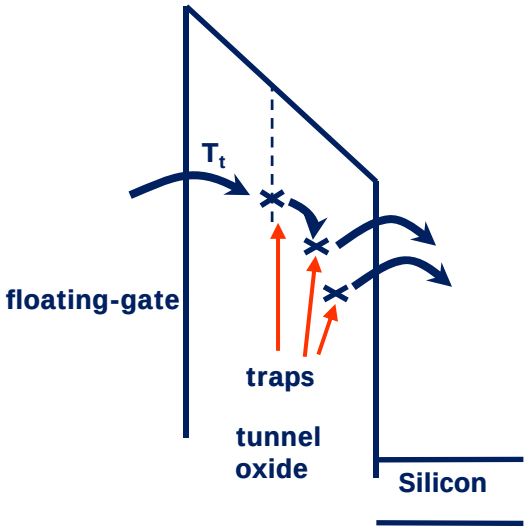

NOR (ETOX & NROM) Scaling:
3.2ev required to surmount Si-SiO2 barrier → Limits Cell Gate Length Scaling

NOR Reliability: Write / Erase Tox traps leading to TAT or de-trapping of trapped oxide electrons

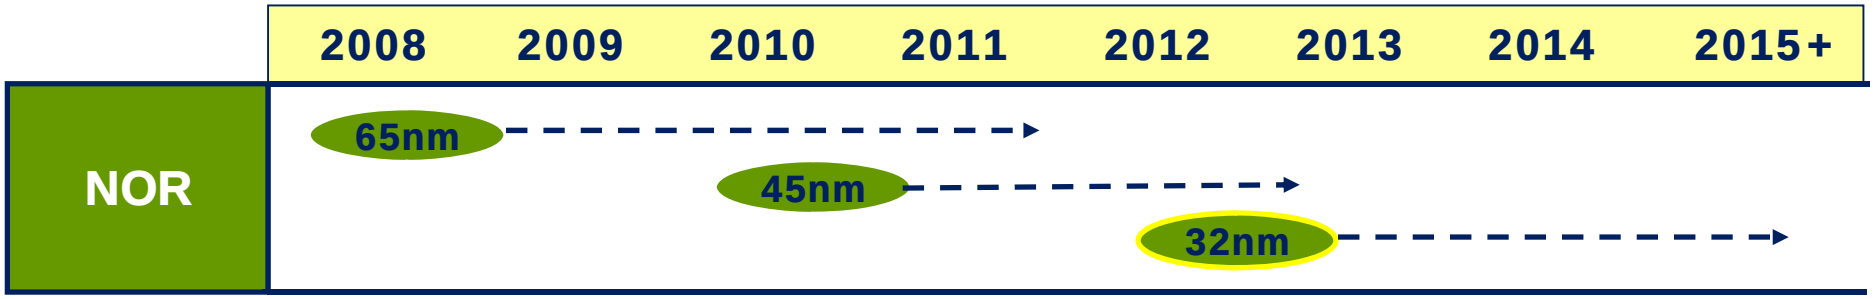

NAND Scaling

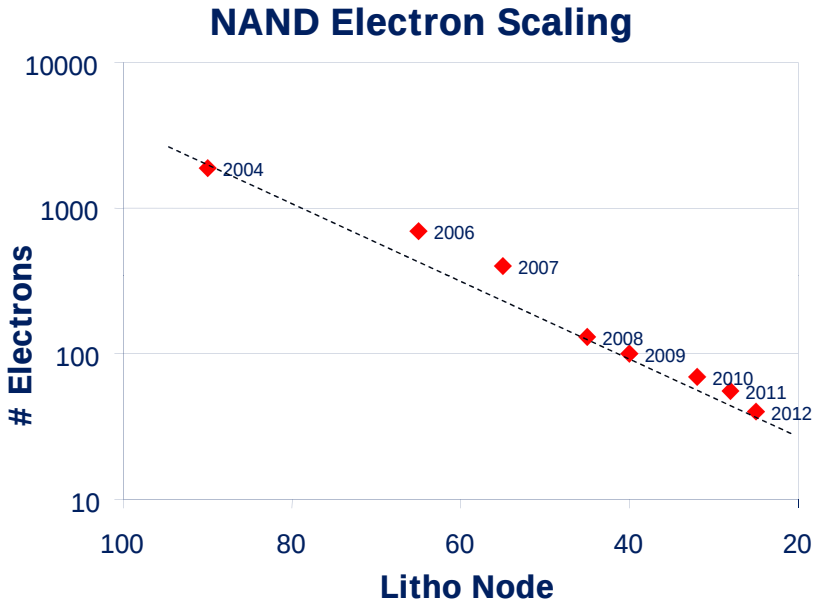

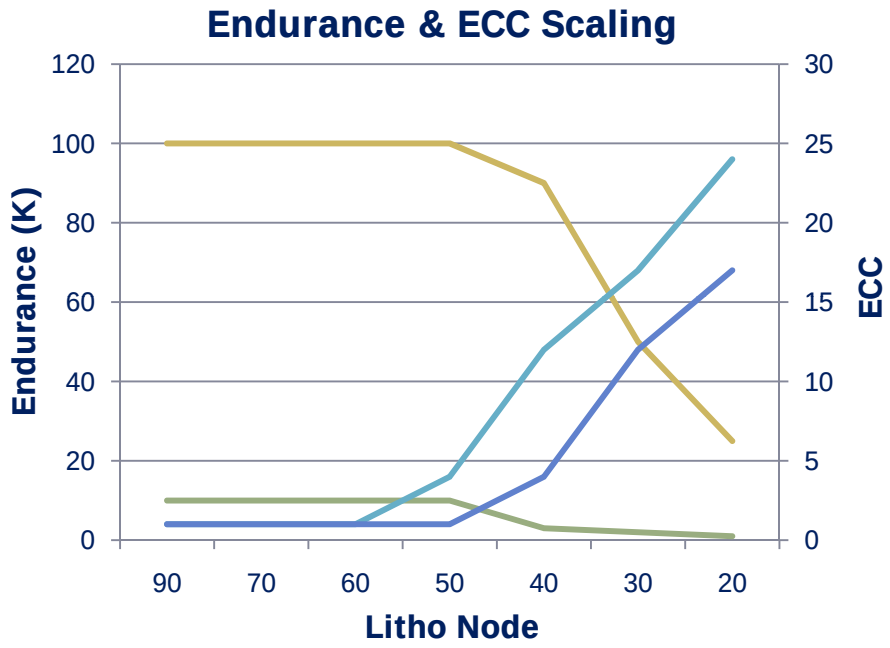

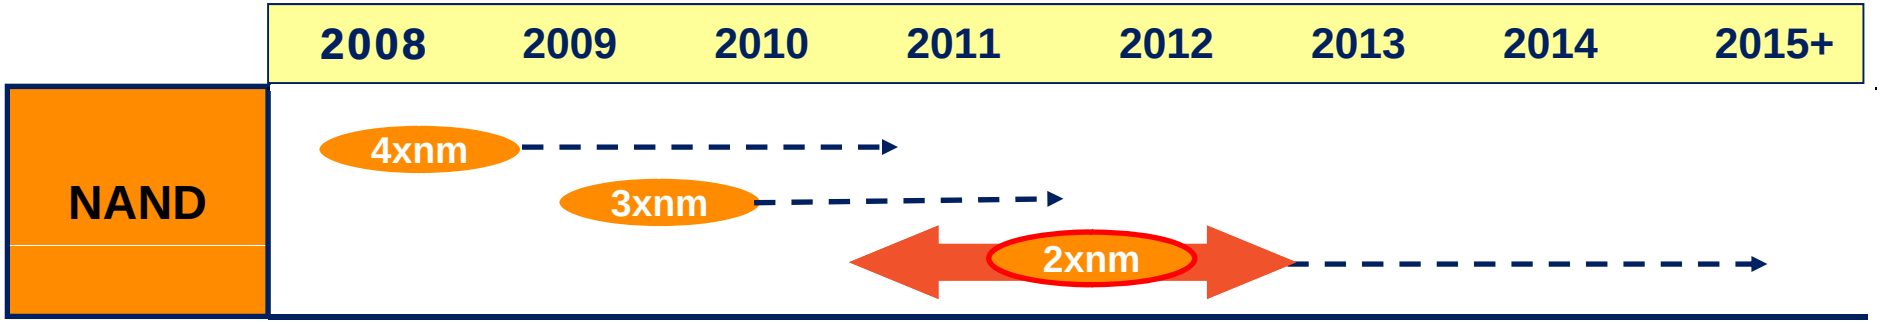

The 10's

NOR / NAND / DRAM

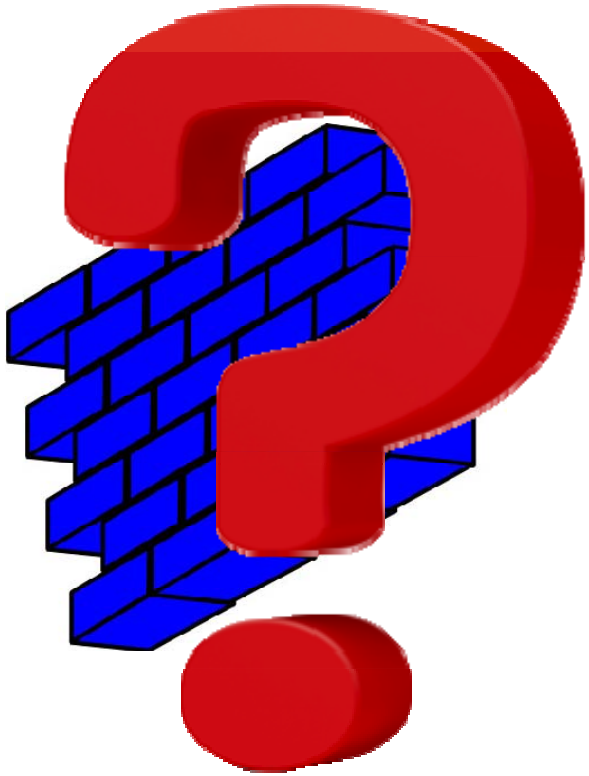

- The technology "Brick Wall"
- What can /should we do about it?
- What's beyond the brick wall?

Phase Change Memory

- Storage
  - GST: Germanium-Antimony-Tellurium Chalcogenide glass
  - Cell states varying from amorphous (high resistance) to crystalline (low resistance) states
- Read Operation
  - Measure resistance of the GST
- Write Operation
  - Heat GST via current flow (Joule effect)
  - Time at critical temperature determines cell state

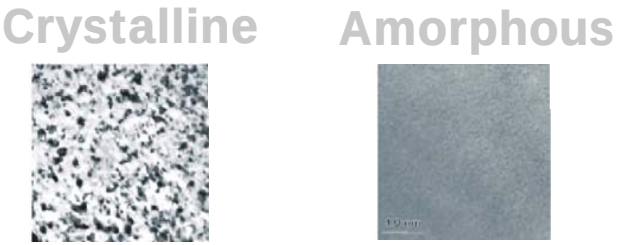

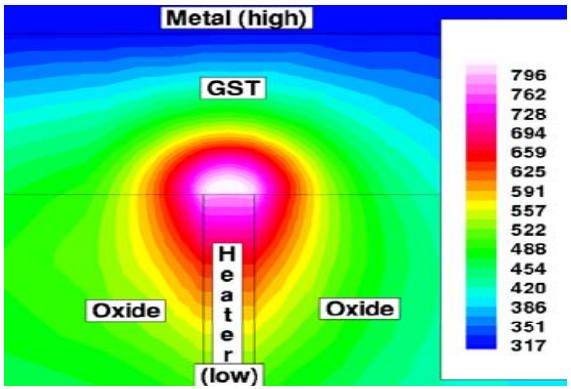

Temperature during write

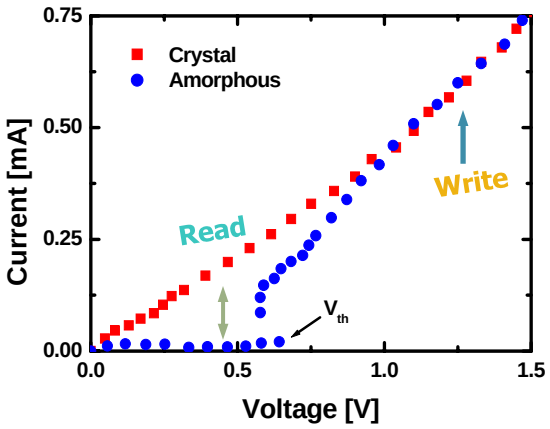

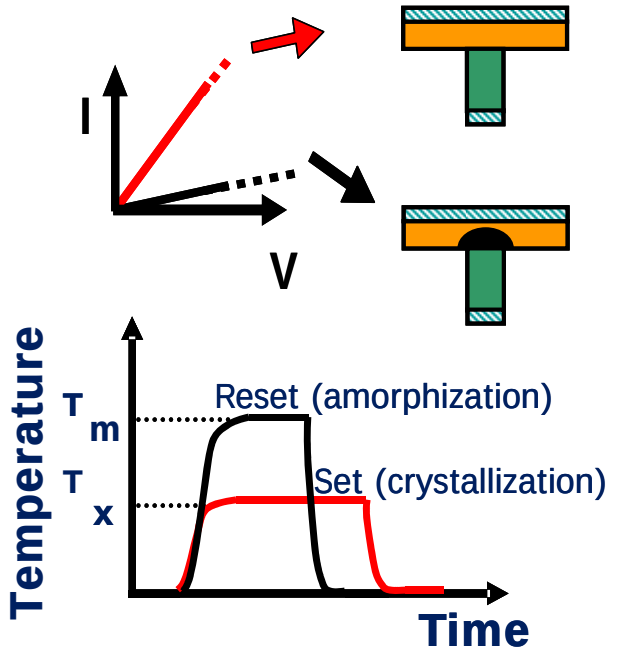

Memory Characteristics
PCM Offers Attributes of RAM & NAND

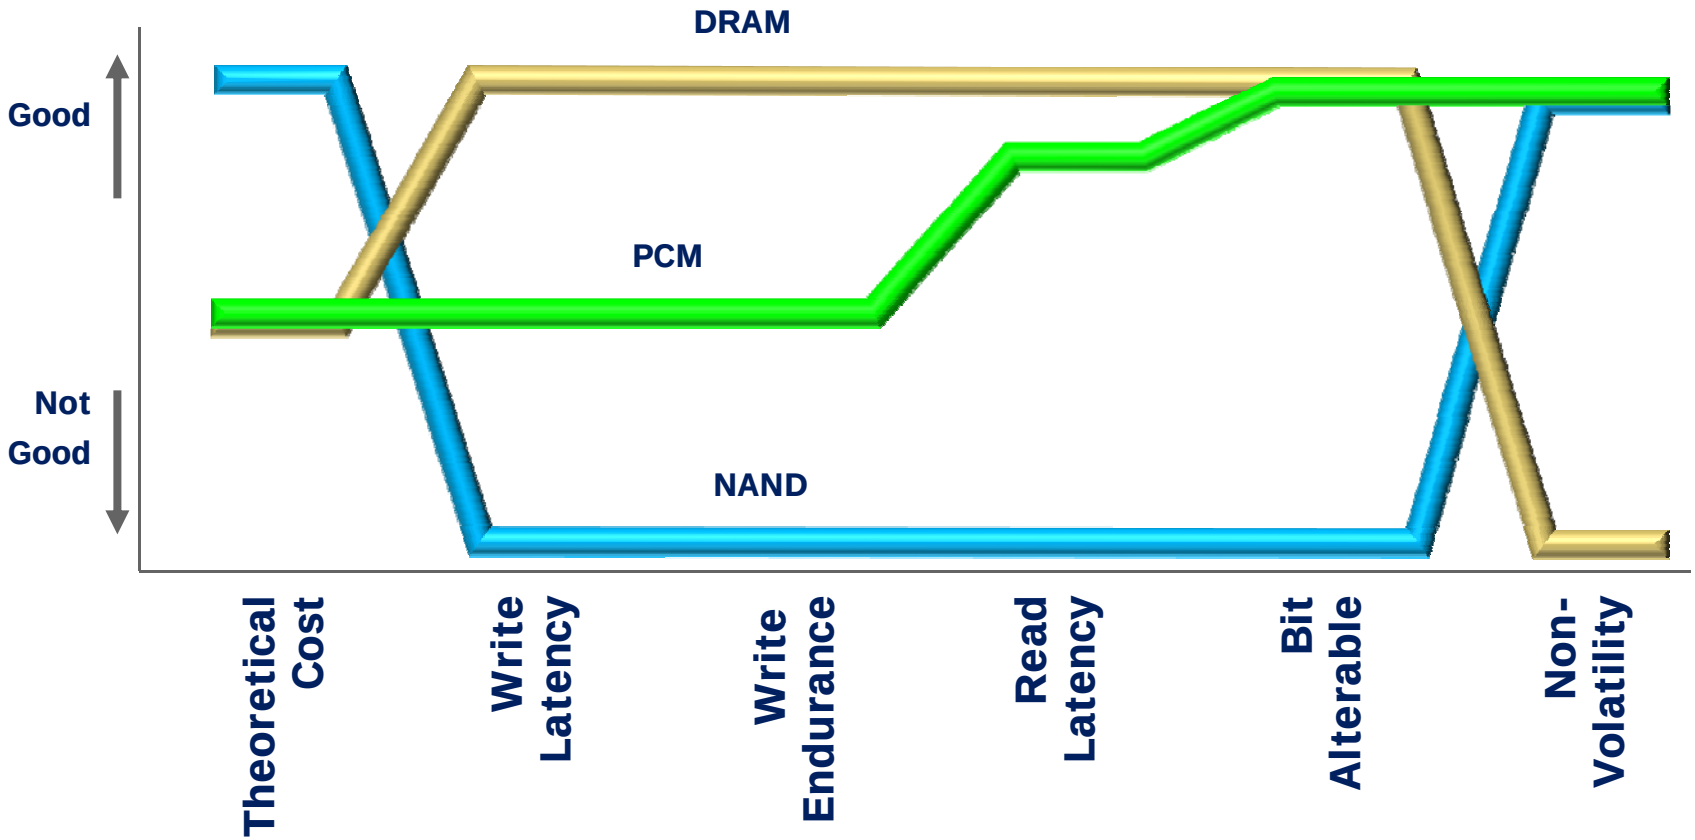

Performance & Density Comparisons

Circa 2011, 45nm Silicon

Attributes	DRAM	PCM	NAND	HDD
Non-Volatile	No	Yes	Yes	Yes
Idle Power	~100mW/GB	~1 mW/GB	~1 mW/GB	~10W
Erase / Page Size	No / 64Byte	No / 32Byte	Yes / 256KB	No / 512Byte
Write Bandwidth	~GB/s per die	50-100 MB/s per die	5-40 MB/s per die	~200MB/s per drive
Page Write Latency	20-50 ns	~1 us	~500 us	~5 ms
Page Read Latency	20-50 ns	~70 ns	~25 µs	~5 ms
Endurance	∞	$10^6 \rightarrow 10^7$	$10^5 \rightarrow 10^4$	∞
Maximum Density	4Gbit	4Gbit	64Gbit	2TByte

Theoretical Cost

Theoretical Chip Cost Factors

Silicon Cost Component		SLC PCM	DRAM	SLC NAND
Die Size	Cell Size (F^2)	5.5	6.0	5.0
	4G Prod Example	1.0x	1.2x	1.0x
Wafer Complexity	Total Process Mask Count	~35	~34	30
	300mm cost structure	1.2x	1.2x	1.0x
Theoretical Die Cost Summary		1.2x	1.4x	1.0x

- PCM will be cheaper than DRAM at lithography parity
- PCM scales to lower densities better than NAND
- PCM attributes can also save cost at system level

Bit Alterability

Bit Alterability
Ridiculously Simple

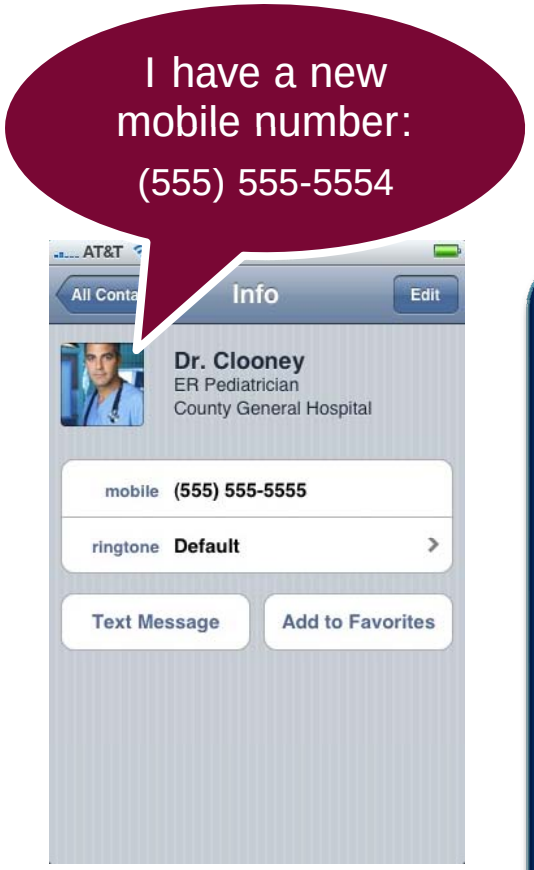

NAND

1. Read 4KB from NAND w/ECC
2. Write to RAM
3. Modify RAM
4. Locate new NAND page
5. Write new NAND page
6. Calculate & Write ECC
7. Mark old NAND page “dirty”
8. Eventually erase NAND block

PCM

1. Write 1 bit in PCM

Much less bus traffic

“Hidden” Power

Ridiculously Simple

Endurance & Retention

Reliability
System Implications

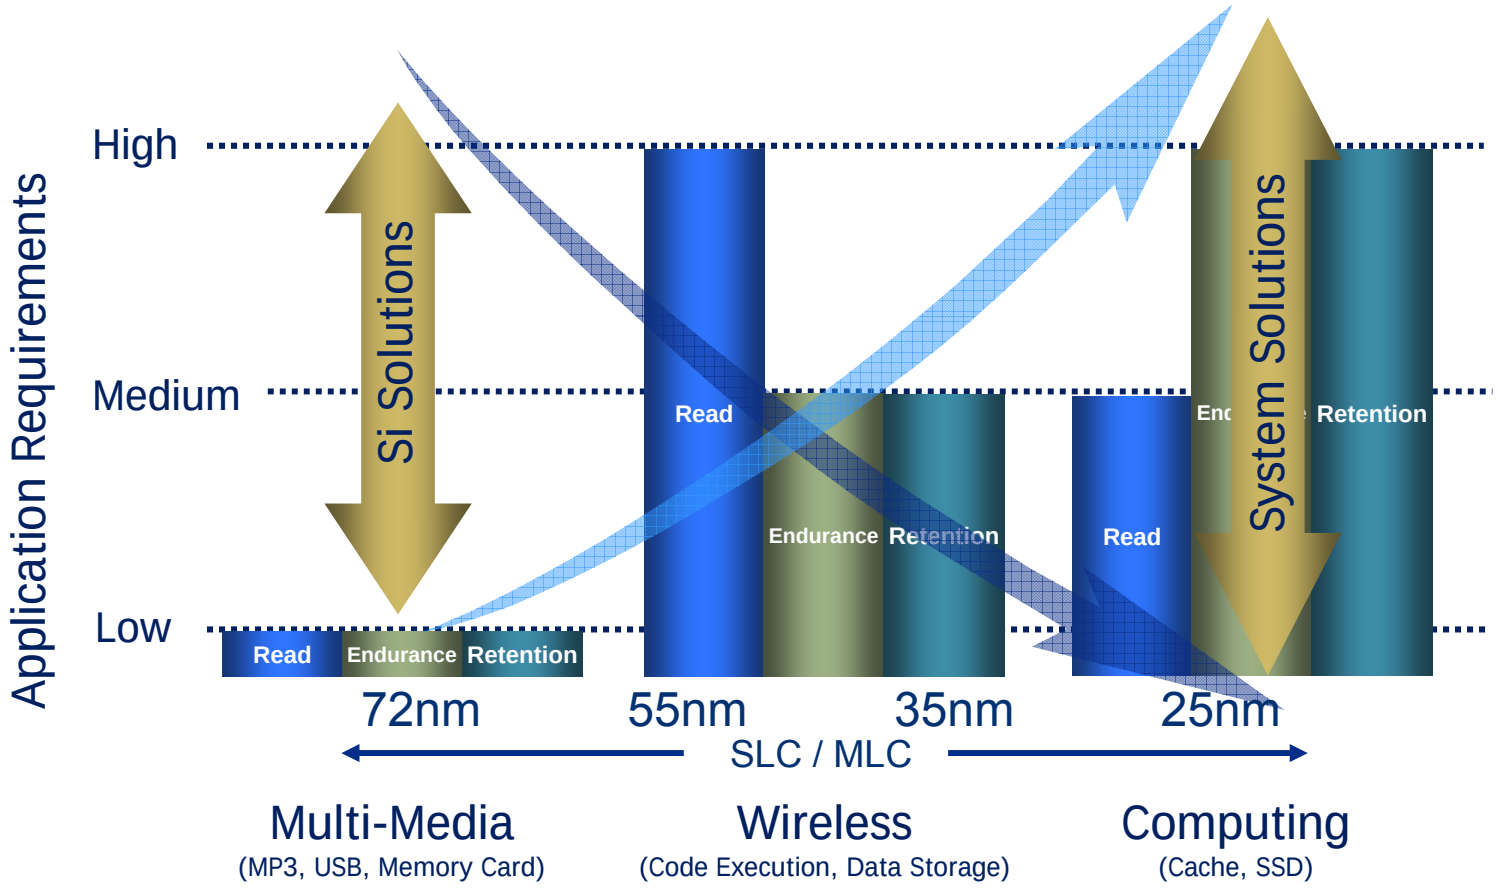

Endurance
Scalability

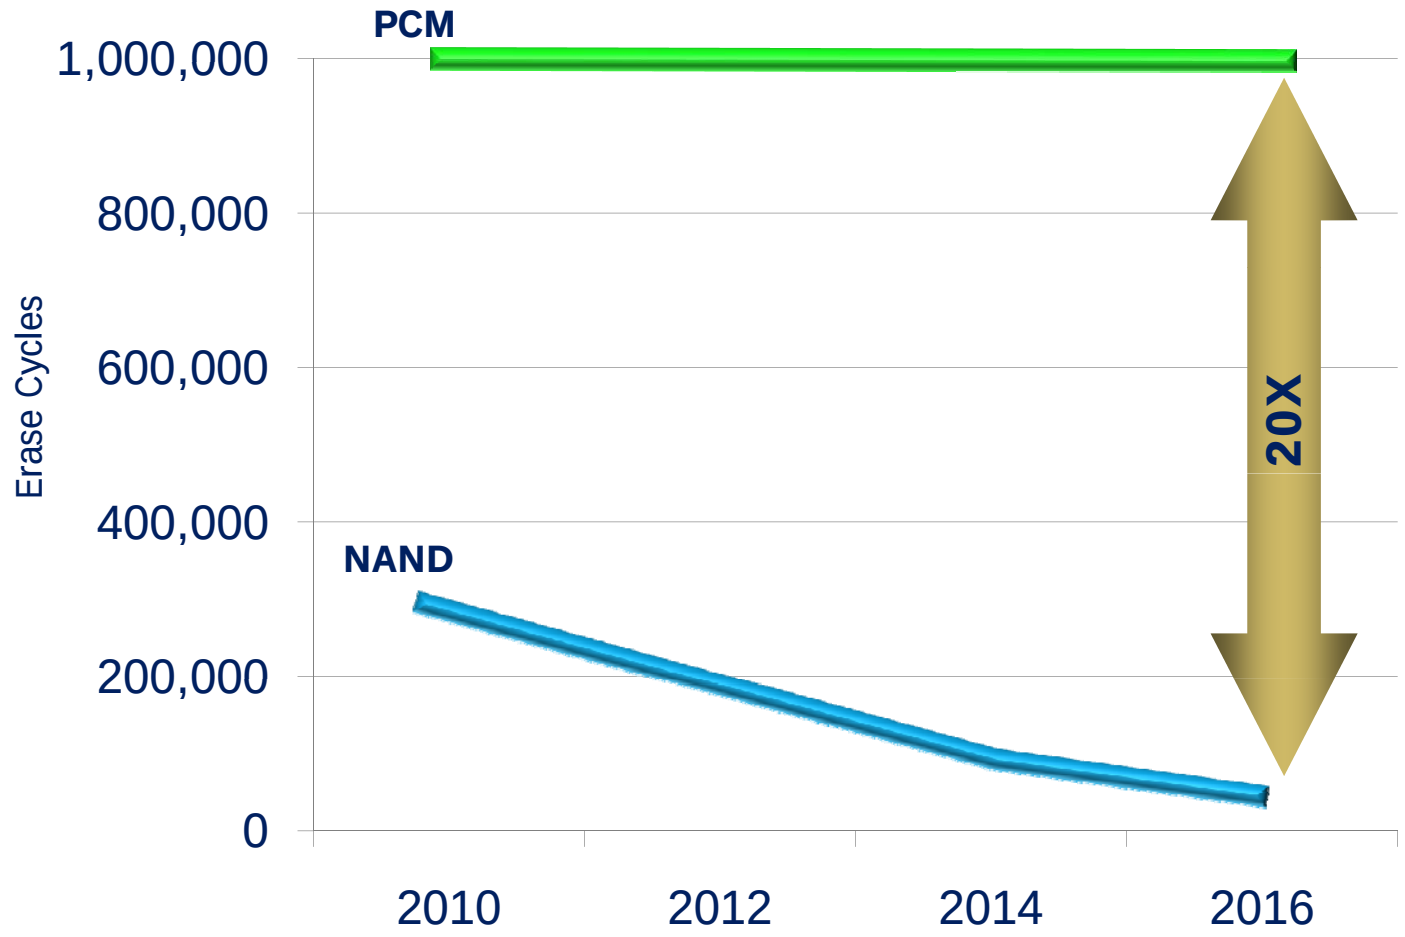

Source: My Estimates

Data Retention
Historical View

200k year?

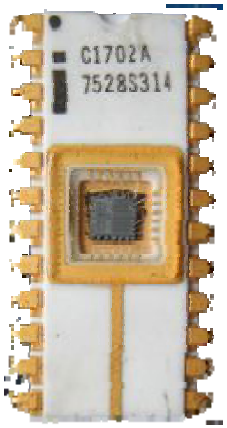

IEEE 1005-1998
(Revision of IEEE Std 1005-1991)

IEEE Standard Definitions and Characterization of Floating Gate Semiconductor Arrays

Sponsor
Standards Committee of the IEEE Electron Devices Society

Approved 25 June 1998
IEEE-SA Standards Board

"Endurance is a measure of the ability of a nonvolatile memory device to meet its data sheet specifications as a function of accumulated nonvolatile data rewrites or program/erase cycles."

JEDEC STANDARD

Electrically Erasable Programmable ROM (EEPROM) Program/Erase Endurance and Data Retention Stress Test

JESD22-A117A
(revision of JESD22A117, January 2000)

MARCH 2006

JEDEC Solid State Technology Association

JEDEC EIA

10 yr at "low cycles"

1 yr at data sheet

200 ka — **1971** — **1998** — **2006**

News Flash: PCM Retention is NOT a function of endurance

Endurance
SSD: Fast Growing Segment

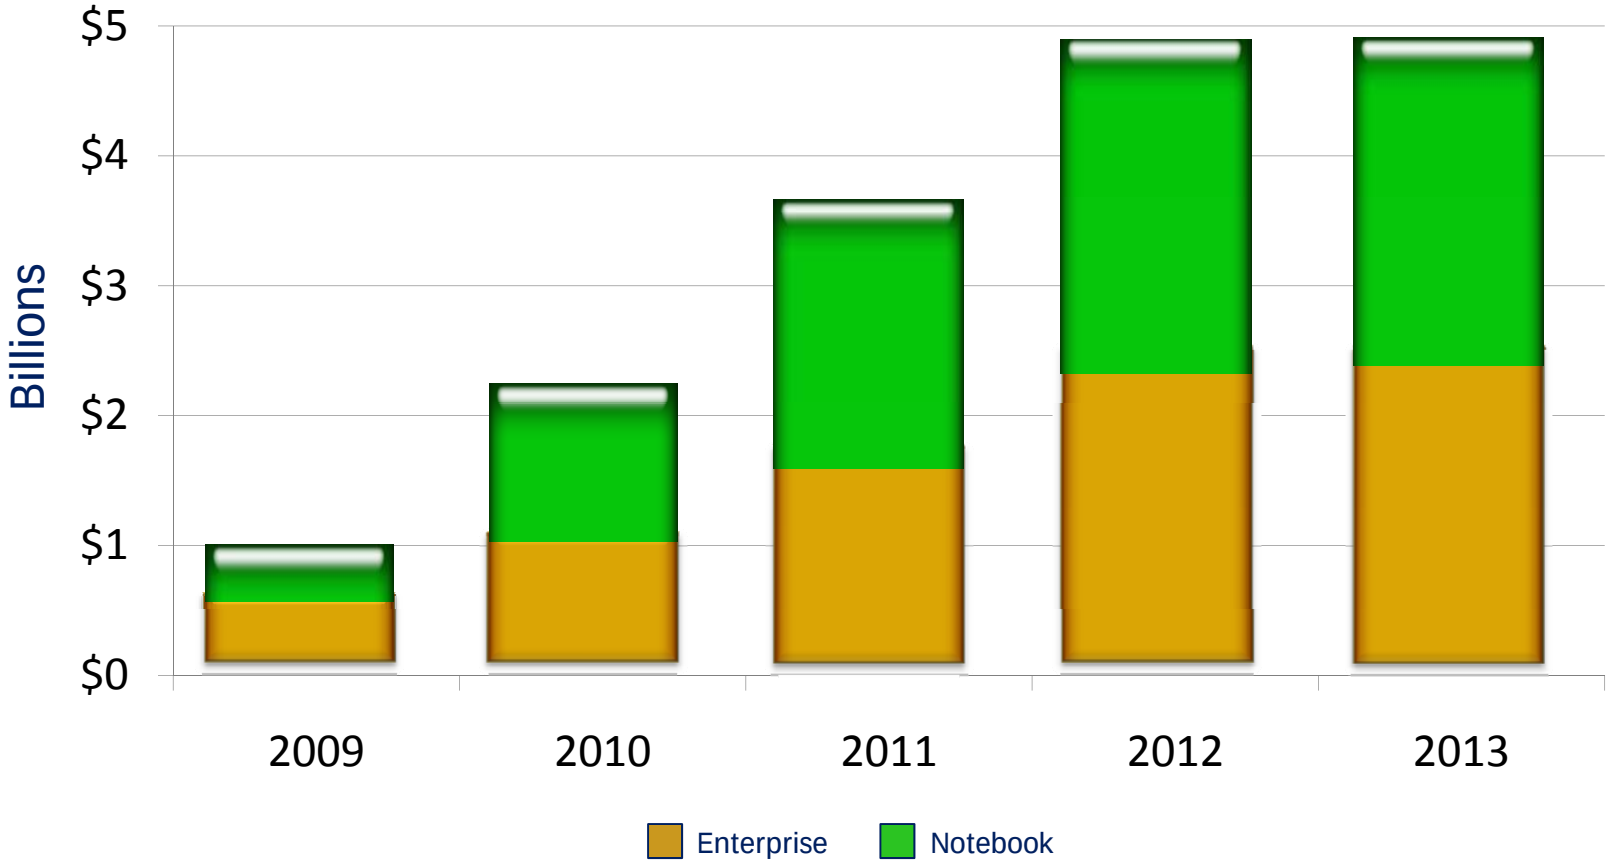

Source: iSupply, Dec 2009

Interfaces
Getting Faster

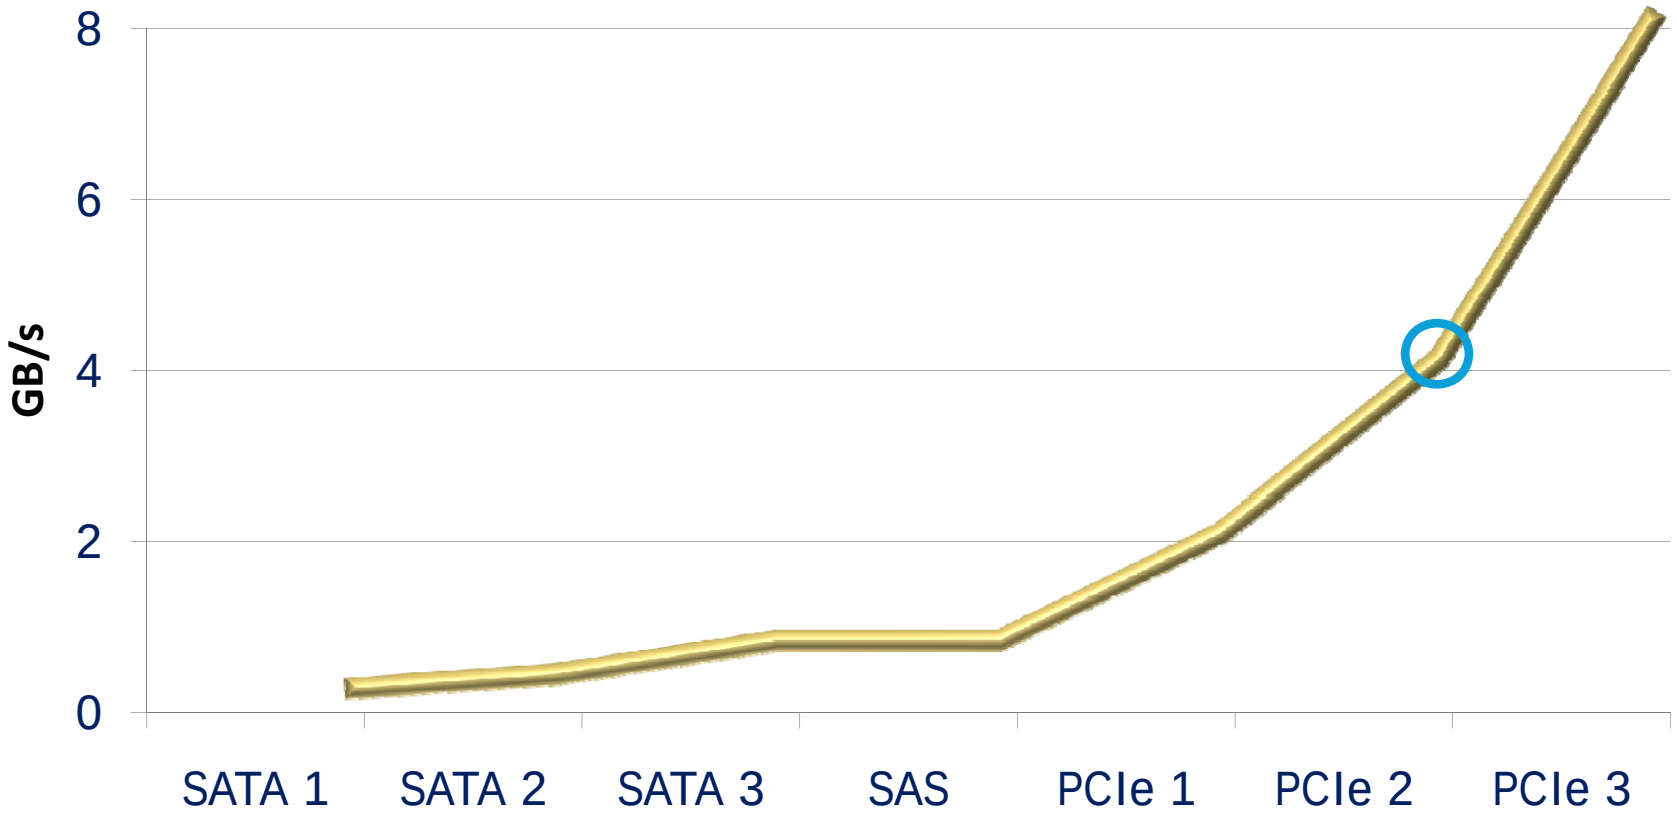

4GB/s at 90/10 Read/Write Equates to 34TB's per day

Note: PCIe x8

System Solutions
Endurance vs. Density

PCIe2

90/10

5 yrs

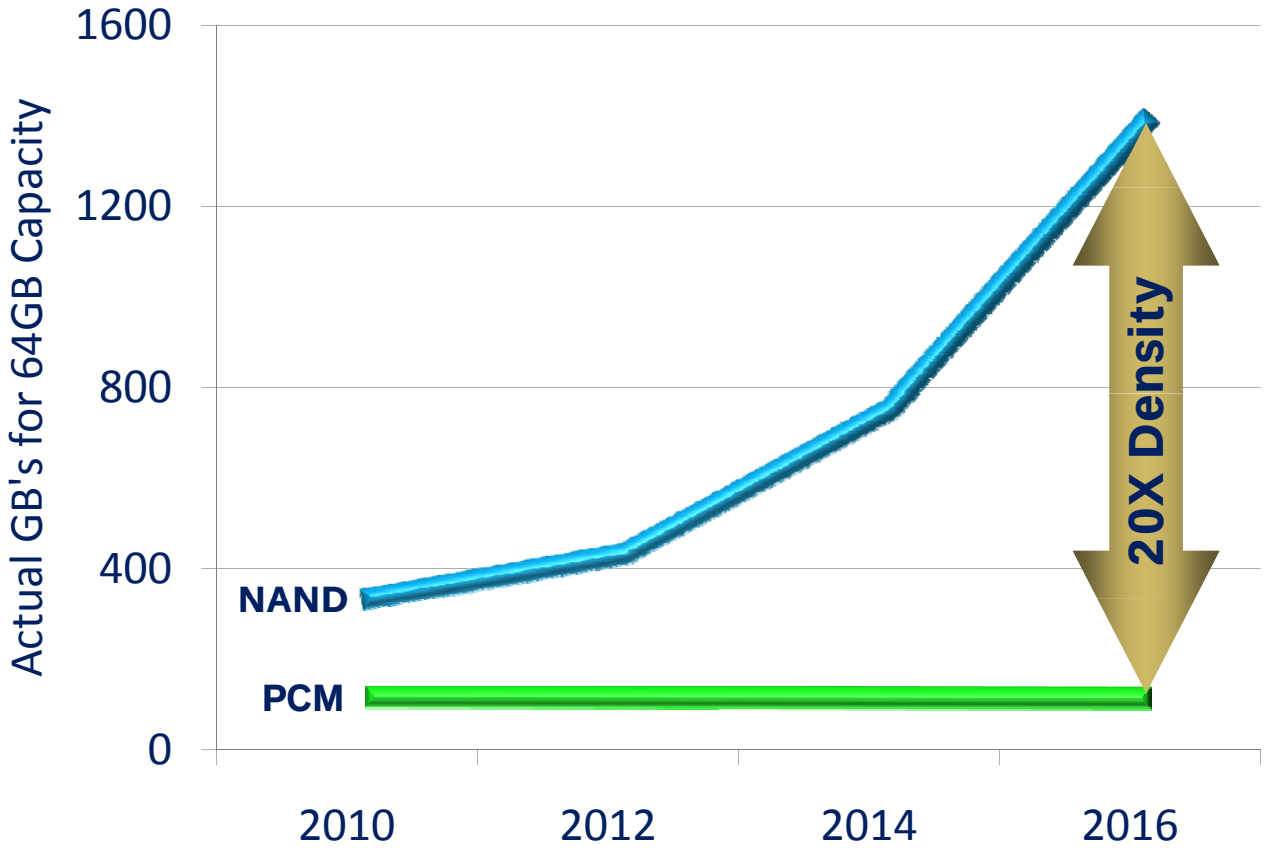

Latency

Performance
Bandwidth vs. Latency

Bandwidth

"Add More"

0-60mph → 10.4s
$20,000

Latency

"Add Something Faster"

0-60mph → 3.5s
$200,000

Application to Compute Approach is Changing

Bandwidth Improvement

Latency Improvement

or

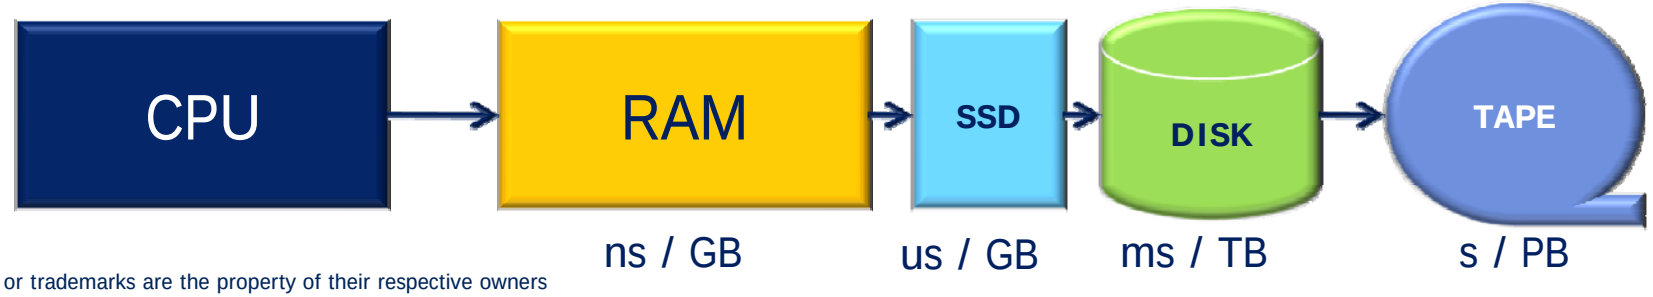

Performance
Relatively Speaking

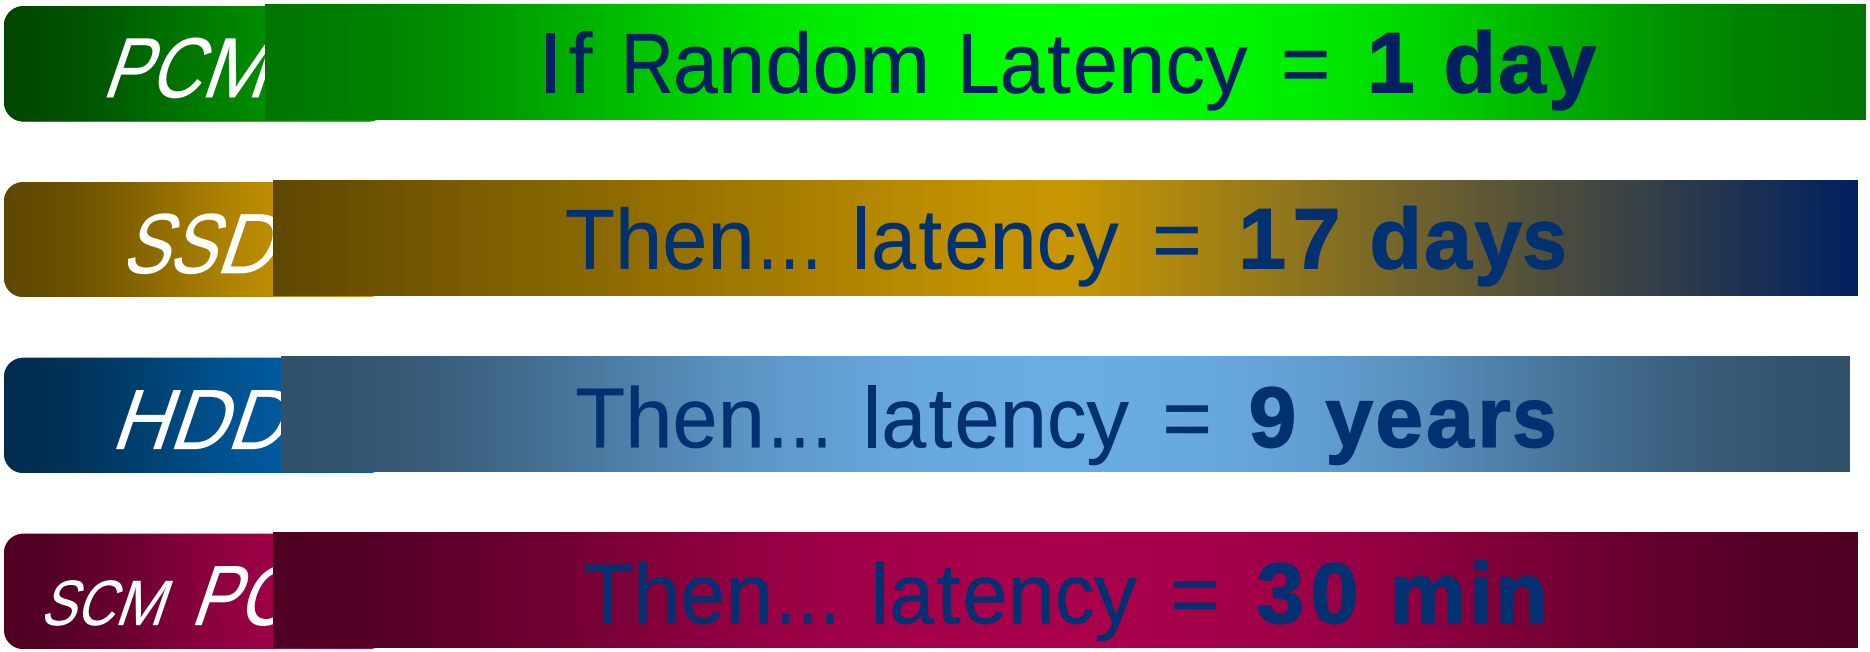

Storage Class PCM can perform >800 data base searches in the time it takes an SSD to do 1

Latency
Value goes up with information growth

amazon.com "… every 100ms of latency cost them 1% in sales"

Google™ "…an extra 500ms in search page generation time dropped traffic by 20%"

TABB GROUP "… a broker could lose $4 Million Dollars per millisecond if their electronic trading platform is 5ms behind the competition"

Source: http://highscalability.com

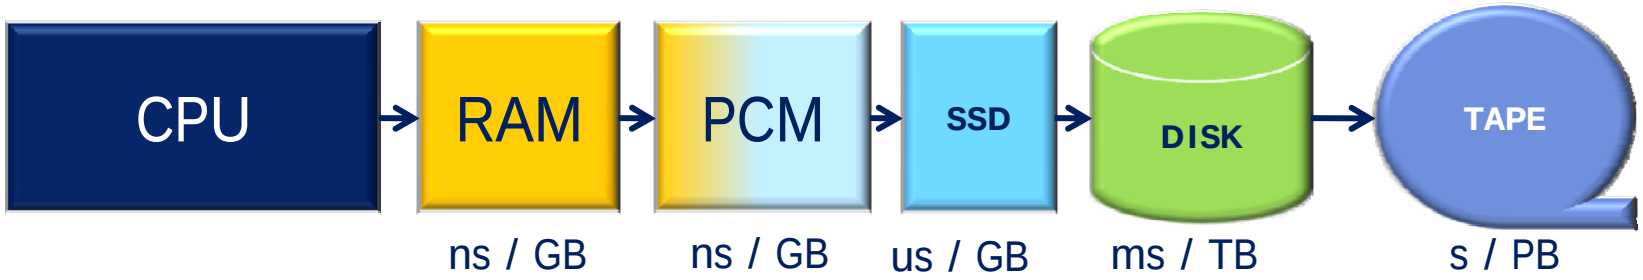

Latency
Still Not Convinced?

High Frequency Trader: Trading In Mere "Millisconds" Is Like, So Last Century

Courtney Comstock | Mar. 4, 2010, 10:07 AM | 852 | 5 | Print

Tags: Wall Street, High Frequency Trading

When we made the mistake of saying "milliseconds," instead of "microseconds" at a Mankoff Group HFT conference the other night, we were quickly corrected.

"Trading in milliseconds is from like, last century," a trader told us.

"The new frontier is microseconds," another added.

"We're trying to get to a speed where we can trade in 5 microseconds - actually it might even be down to 3 now."

PCM
Industry Demands Will Drive Change

Bit Alterable

Ridiculously Simple

Endurance

High Reliability Usage Models

Latency

Instant Access

Retention

Not a function Of Endurance


Micron